Answer Engine Optimization in 2026

Mechanisms, Evidence, and Practice in the Age of AI-Native Search

A Research Review - Marlin Media

July 2026

Abstract

Search behavior shifted materially between 2024 and 2026 as large language model (LLM)-powered assistants and AI-generated summaries became a primary interface for information discovery.

This paper reviews the discipline that has emerged in response — Answer Engine Optimization (AEO) — and situates it alongside two adjacent terms, Generative Engine Optimization (GEO) and Large Language Model Optimization (LLMO).

Drawing on the foundational Princeton/Georgia Tech generative-engine study and its 2025-2026 replications, together with current industry measurement reports, this review synthesizes what is empirically known about how answer engines select, weight, and cite sources; how citation behavior differs across Google AI Overviews, ChatGPT, Perplexity, Microsoft Copilot, Gemini, and Claude; and what content and technical practices show a repeatable, measurable relationship to citation likelihood. The review closes with a practical framework, a summary of open methodological questions, and a caution against vendors promising guaranteed AI placement, which the underlying research does not support.

1. Introduction

By mid-2026, a growing share of consumer and business research begins not with a list of ranked links but with a synthesized answer.

Analysts at Gartner projected that roughly a quarter of organic search traffic would migrate from traditional result pages to AI chatbots and virtual assistants by 2026, and independent tracking of Google's rollout shows AI Overviews now appearing on a large share of queries, with measurable reductions in click-through to the underlying sources.

At the same time, dedicated AI assistants have themselves become large-scale answer surfaces: OpenAI's ChatGPT alone is now handling billions of queries per week, and a wide range of vertical and horizontal competitors — Perplexity, Microsoft

Copilot, Google's Gemini and AI Mode, and Anthropic's Claude — have grown alongside it.

The scale of this shift is now well documented. OpenAI reported that ChatGPT reached 900 million weekly active users in February 2026, more than double the 400 million reported a year earlier, with the platform processing on the order of 2.5 billion prompts per day.

Google's AI Overviews reach roughly two billion monthly users across more than 200 countries, and independent tracking by BrightEdge placed AI Overviews on about 48 percent of monitored queries by February 2026, a 58 percent year-on-year increase, with health, education, and research verticals triggering them on roughly 80 percent of queries.

Perplexity, the largest dedicated citation-first answer engine, passed 45 million monthly active users and roughly 780 million monthly queries during 2025, more than doubling its user base within a year. Against this backdrop, Gartner's early-2024 prediction that traditional search engine volume would fall 25 percent by 2026, initially treated as aggressive, now reads as directionally correct and possibly conservative for informational verticals.

This shift changes what "visibility" means for a publisher, brand, or content team. In a link-based search result, ranking in the top positions is both necessary and largely sufficient for visibility.

In an answer-generation pipeline, a page must additionally be retrieved as a candidate, judged trustworthy and relevant relative to a small number of competing sources, and then either quoted, paraphrased, or named as a citation inside a synthesized answer. Answer Engine Optimization is the set of practices aimed at improving a page's odds at each of those stages.

This paper does not promote any single vendor's tool or methodology. It synthesizes publicly available research and cross-checks recurring claims across multiple independent sources, flagging where the evidence is strong (replicated, peer-reviewed, or drawn from large samples) and where it is closer to industry anecdote.

2. Defining the Terms: SEO, AEO, and GEO

Industry usage of "AEO" and "GEO" is inconsistent, and the two terms are frequently used interchangeably. The distinction that recurs most often across 2026 sources is one of scope and layer, not of competing methodology.

2.1 Search Engine Optimization (SEO)

SEO remains the discipline of earning visibility in a ranked list of links, optimizing for crawlability, indexation, backlink authority, and keyword relevance. Most sources agree SEO remains the foundation on which AEO and GEO are built: a page an answer engine cannot find or trust as a web source is rarely eligible to be retrieved at all.

2.2 Answer Engine Optimization (AEO)

AEO is the narrower, answer-layer discipline: structuring a specific page or passage so that a system can extract a clear, self-contained answer to a specific question and, where relevant, attribute it. Typical tactics include leading with a direct answer, using question-formatted headings, applying FAQ, HowTo, and Article schema markup, and keeping each section independently understandable outside the context of the rest of the page.

2.3 Generative Engine Optimization (GEO)

GEO is generally described as the broader, brand-level discipline: shaping how a brand is represented across the entire source ecosystem an LLM draws on — owned content, earned media, review platforms, forums, and structured data — so that when a model synthesizes an answer, it both retrieves the brand's content and frames it favorably. Where AEO targets a single answerable moment, GEO targets the aggregate impression a model has formed about an entity across its training and retrieval sources.

2.4 Large Language Model Optimization (LLMO)

A smaller number of 2026 sources use "LLMO" as a near-synonym for GEO, emphasizing optimization for the model's internal weighting of entities and facts rather than any single platform's retrieval mechanics. In practice, the three terms describe overlapping work, and most practitioner guides in 2026 treat them as a continuum rather than three separate workstreams.

2.5 Comparative Summary

Primary goal	Rank in a list of links	Be extracted as a direct answer	Be cited and framed favorably in a synthesized answer
Unit optimized	Page / site	Passage / section	Entity / brand across the source ecosystem
Core levers	Keywords, backlinks, technical SEO	Structure, schema, direct-answer formatting	Statistics, citations, quotes, consistent entity data, earned media
Success metric	Rankings, impressions, clicks	Featured snippets, extraction rate	Citation frequency, share of voice across platforms

Table 1. A composite comparison of SEO, AEO, and GEO as described across 2026 industry sources.

3. How Answer Engines Select and Cite Sources

Most consumer-facing answer engines rely on some form of retrieval-augmented generation (RAG): rather than answering solely from parametric knowledge learned during training, the system issues a live search, retrieves candidate documents, and conditions its generated answer on that retrieved content. Understanding this pipeline clarifies where optimization can and cannot help.

3.1 Query interpretation

The engine parses the user's question into an intent and a set of underlying entities and relationships rather than a literal keyword string. This is why a

page can be retrieved for a query that never uses its exact phrasing, provided the underlying concepts are covered clearly.

3.2 Retrieval

The system searches an index — frequently a conventional web index licensed from or built like Google's or Bing's — for documents that are semantically close to the parsed query, producing a shortlist of candidates rather than the full ranked result set a user would see in classic search.

3.3 Ranking and selection

Candidates are scored on relevance, apparent authority, recency, and structural quality — whether the page states a clear answer, backs it with evidence, and is organized so key facts can be isolated from surrounding content.

3.4 Synthesis and citation attribution

The model reconciles overlapping or conflicting facts across the surviving candidates and drafts a single answer. Sources that contributed a specific, extractable fact typically receive an inline citation; sources that contributed nothing usable, or whose claims were contradicted by a stronger source, are dropped even if they were part of the original retrieval set. One recurring finding across 2026 studies is that engines increasingly cite passages rather than whole pages, which shifts the optimization target from an entire article to each individual paragraph's ability to stand alone.

4. The Empirical Evidence Base

4.1 The founding study: Aggarwal et al. (Princeton, Georgia Tech, IIT Delhi; KDD 2024)

The most frequently cited empirical foundation for GEO is a large-scale study presented at the 2024 ACM SIGKDD conference, which tested a set of content

modifications across roughly ten thousand queries and measured their effect on whether a generative answer engine cited a given source. The study identified several tactics with a consistent, positive relationship to citation probability: adding statistics and data points, adding direct quotations (often attributed to a named expert), adding citations to external authoritative sources, and rewriting passages for greater fluency and readability. Reported effect sizes in the original study and its replications cluster in the 25 to 41 percent relative-visibility range for the strongest individual tactics, with citation-adding and statistics-adding showing the largest and most consistently replicated gains. Combining fluency optimization with statistics addition produced a larger effect than either tactic alone.

A methodologically important secondary finding is a leveling effect: sources that started in a weak competitive position — effectively a low-ranking page in a five-source shortlist — gained disproportionately more from optimization than already-strong sources did, with some analyses reporting visibility gains above 100 percent for the weakest starting position. This suggests GEO-style optimization may lower, rather than reinforce, the advantage of large incumbent sites in a given answer, though this claim comes from a controlled experimental setting and has not been confirmed at the same magnitude in live production systems.

The study also produced a negative finding worth emphasizing for anyone importing classic SEO habits into this space: keyword stuffing measurably reduced citation likelihood in at least one tested platform, likely because it degrades the fluency signal models use as a proxy for source quality.

4.2 Follow-up and replication research

Since 2024, several follow-up studies have tested the original findings under more realistic conditions. A 2025 benchmark of conversational-SEO tactics, evaluated across a wider set of prompts and platforms, found that many popularly recommended tactics show little or no measurable effect once tested outside a controlled, small-competitor-set environment, and that plain topical relevance of the source remains the strongest single predictor of

citation. This is a useful corrective to marketing claims that treat the original 2024 effect sizes as guaranteed real-world lifts: the original study used a zero-sum setup with only five competing sources per query, which tends to inflate relative gains, and its content modifications were generated by an LLM rather than applied by human editors, leaving open how well the findings transfer to normal editorial workflows.

Industry-run studies analyzing citation logs at scale add further texture. One large analysis of citation patterns found that only about a tenth of sites were cited by two major assistants simultaneously, underscoring that different engines draw on meaningfully different source pools rather than one shared "AI index." Separate research tracking overlap between classic organic rankings and AI citations found that overlap has fallen sharply — from roughly three-quarters overlap in mid-2025 to well under half in early 2026 on some platforms — meaning a page's position in traditional search results is a weaker and weaker proxy for its odds of being cited in an AI answer.

Source-concentration analyses sharpen this picture further. Profound's longitudinal dataset of more than 680 million tracked citations across ChatGPT, Google AI Overviews, and Perplexity found that Wikipedia accounts for roughly 47.9 percent of ChatGPT's top-ten source share while Reddit accounts for roughly 46.7 percent of Perplexity's, and that only about 11 percent of domains are cited by both platforms. Semrush's 13-week study of 230,000 prompts documented how volatile these patterns are: ChatGPT cited Reddit in close to 60 percent of responses in early August 2025 before collapsing to around 10 percent by mid-September, coinciding with Google's removal of the num=100 results parameter, while citation shares for the same domains held broadly steady on AI Mode and Perplexity. A separate 30-day analysis of 17,551 citations within a single B2B category found vendor-owned content accounted for under 1 percent of citations, with third-party editorial, community, and reference sources dominating, consistent with broader industry estimates that roughly 94 percent of AI citations trace to earned rather than owned media.

4.3 Freshness and structural findings

Multiple 2026 industry reports converge on freshness as a strong, controllable predictor, particularly for commercial and evaluation-stage queries: one large-scale citation analysis found that the large majority of citations for buyer-intent queries came from pages updated within the previous twelve months, with more than half refreshed within the previous six. A separate structural finding, repeated across several 2026 guides, is that a large share of citation-worthy passages come from early in a page — content in roughly the first third of an article accounts for a disproportionate share of what gets quoted or cited, favoring a direct-answer-first structure over a build-up-to-the-point narrative style.

The freshness effect has since been quantified with more precision: AirOps' 2026 citation analysis found that 95 percent of ChatGPT citations came from content updated within the previous ten months, and that pages displaying a visible "last updated" timestamp earned roughly 1.8 times more citations than comparable pages without one.

4.4 Summary of Reported Effects

Adding statistics / data points	≈ +22% to +41%	Strong — replicated across original study and multiple follow-ups
Adding direct, attributed quotations	≈ +28% to +41%	Strong — one of the two largest effects in the original study
Adding citations to authoritative sources	≈ +28% to +115% (largest for low-starting-position pages)	Strong, but magnitude highly context-dependent
Fluency / readability rewriting	Positive; largest when combined with statistics	Moderate — consistent direction, variable size
Content freshness (updated within 6–12 months)	Associated with large majority of citations on commercial queries	Moderate — correlational, from large industry samples

Keyword stuffing	≈ -9% to -10% (negative)	Moderate — direction consistent, one primary source
Generic "conversational SEO" tricks (broad category)	Little to no measurable effect once tested at scale	Emerging — from 2025 benchmark work correcting early hype

Table 2. Reported directional effects of common GEO/AEO tactics, synthesized across the Princeton/Georgia Tech study and its 2025–2026 follow-ups. Figures are relative visibility changes reported by the cited studies, not guarantees of real-world performance.

5. Platform-Specific Citation Behavior

Because each answer engine is built on a different retrieval stack, licensing arrangement, and model, they do not behave identically even when given the same query. 2026 industry guides converge on a set of rough platform tendencies, summarized in Table 3. These should be read as directional patterns reported by third-party analyses rather than confirmed engineering specifications, since none of the underlying platforms publish their full ranking logic.

Google AI Overviews / AI Mode	Draws heavily from pages already ranking in the top organic results	Traditional SEO fundamentals remain a prerequisite, not an alternative
ChatGPT Search	High query volume; increasingly cites passages rather than whole pages; strong weighting toward freshness and structured, information-dense content	Prioritize direct-answer formatting and frequent content refreshes
Perplexity	Rewards freshness, source authority, and presence across multiple channels simultaneously	Maintain consistent facts and entity data across owned site, listings, and third-party profiles
Microsoft Copilot	Leans heavily on LinkedIn and professional-network content for B2B-	For B2B brands, a maintained LinkedIn presence functions as

	style queries	an AEO channel, not just a social one
Gemini	Analyzes multimodal content, including video and image context, alongside text	Video transcripts, captions, and structured image metadata are citation-relevant, not cosmetic
Claude	Reported preference for comprehensive, long-form, well-sourced explainer content over short snippets	Depth and topical completeness may matter more than snippet-style brevity

Table 3. Reported platform-specific tendencies, synthesized from multiple 2026 industry analyses. Not independently verified against platform source code.

6. Market-Level Shifts Observed in 2026

Beyond individual-page tactics, three market-level trends recur across 2026 sources and are relevant to how organizations should prioritize AEO investment.

- Declining organic click-through on AI-augmented queries. Independent analyses report organic click-through rates falling sharply — by roughly 60 percent in some samples — on Google queries where an AI Overview is present, consistent with earlier Pew Research findings that AI summaries reduce the likelihood a user clicks through to any source at all.

  The most rigorous figures now available: Ahrefs' updated analysis of 300,000 keywords using aggregated Search Console data found the presence of an AI Overview correlates with a 58 percent lower click-through rate for the top-ranking page as of December 2025, up from a 34.5 percent reduction measured in April 2025; and Pew Research Center's behavioral study of 68,879 real searches by 900 U.S. adults found users clicked a traditional result on only 8 percent of visits when an AI summary was present versus 15 percent without, clicked a link inside the summary just 1 percent of the time, and ended their session entirely 26 percent of the time versus 16 percent without a summary.

  Similarweb data show zero-click searches rising from 56 to 69 percent of queries between May 2024 and May 2025. Two nuances temper the headline

decline: Amsive's 700,000-keyword analysis found branded queries with AI Overviews actually gained click-through (+18.68 percent), and Seer Interactive's study of 25.1 million impressions found that being cited inside the Overview delivered 35 percent more organic clicks and 91 percent more paid clicks than not being cited, making the citation itself the new "position zero."

- Rising AI-referral traffic quality. Despite fewer total clicks, several analyses report that traffic referred from AI assistants converts at a substantially higher rate than average organic traffic, plausibly because a user who clicks after reading a synthesized answer is further along in an evaluation process than a casual searcher landing on a results page.

  Reported multiples vary by methodology but agree on direction: Semrush's cross-industry 2026 data puts AI-referred visitors at roughly 4.4 times the conversion rate of standard organic traffic; Ahrefs' analysis of its own site found AI-referred visitors made up 0.5 percent of sessions but drove 12.1 percent of signups; and Adobe Analytics reported that AI-referred shoppers to U.S. retail sites converted 42 percent better and spent 48 percent more time on product pages than non-AI traffic in March 2026, after underperforming a year earlier.

  Volume remains small, however: Conductor's 2026 benchmarks put AI referrals at about 1.08 percent of total website traffic, with ChatGPT driving roughly 87 percent of visible AI referral visits.
- Falling overlap between organic rankings and AI citations. As noted in Section 4.2, the correlation between a page's traditional search ranking and its likelihood of being cited in an AI answer has weakened significantly since mid-2025, meaning SEO rank and AEO citation success are increasingly separate outcomes that both need to be measured and managed.

Taken together, these trends support treating AI-referred visibility as a distinct KPI — tracked through citation frequency, share of voice on branded and category queries, and AI-assisted conversions — rather than folding it into existing organic-traffic

dashboards, where a genuine gain in AI visibility can be masked by an overall decline in total clicks.

7. A Practical Framework for 2026

Synthesizing the tactics with the strongest and most consistently replicated evidence, four workstreams recur across credible 2026 guidance.

7.1 Structure for extraction

- Lead each page and each major section with a direct, self-contained answer before elaborating.
- Use question-phrased headings that mirror how a user would actually ask the question.
- Keep individual paragraphs short and independently understandable, since engines increasingly cite passages rather than full pages.
- Apply schema markup — FAQPage, HowTo, Article, and Organization/Person for authorship — but only where it accurately reflects visible content; markup that misrepresents the page is a known trust signal to avoid.

7.2 Substantiate every claim

- Back claims with specific statistics and named, verifiable data rather than vague generalizations.
- Use attributed quotations from credible experts where appropriate.
- Cite external authoritative sources inline for factual claims, rather than assuming domain authority alone will suffice.

7.3 Maintain freshness and consistency

- Refresh high-intent, commercial-evaluation content at least every six to twelve months; this is the single most consistently reported correlate of citation on buyer-stage queries.

- Keep factual details — hours, pricing, terms, entity names — consistent across the owned site and third-party listings, review platforms, and social profiles; engines appear to treat cross-source consistency as a trust signal, sometimes referred to informally as a "consensus effect."

7.4 Build off-site authority deliberately

- Invest in earned media and credible third-party coverage; multiple 2026 analyses report that a large majority of AI citations trace back to journalistic and earned-media sources rather than owned content alone.
- Maintain active, substantive presence on the review and community platforms most cited in a given vertical (for example, G2, Trustpilot, or category-specific forums), since domains with an established third-party footprint are reported to have a meaningfully higher chance of being selected as a source.
- Treat structured press releases and public data releases as citation-relevant assets, not just PR formalities.

7.5 Measure the right things

- Track citation frequency and competitive citation share across the platforms relevant to the audience, not just organic rank.
- Track AI-assisted or AI-referred conversions as a distinct funnel, since AI-referred sessions have been reported to convert at rates several times higher than average organic sessions on some platforms.
- Periodically test real queries by hand across major assistants; automated trackers are improving but platform APIs for citation monitoring remain immature and inconsistent across vendors as of mid-2026.

8. Limitations, Risks, and Open Questions

- Effect sizes reported in the founding study come from a controlled, zero-sum experimental design with only five competing sources per query, which tends to inflate relative gains compared to open-web competition; more recent large-

scale benchmarks suggest many popularized tactics show smaller or negligible effects once tested at scale.

- No independent, cross-platform benchmark yet ties specific on-page changes to a guaranteed citation outcome on a live, uncontrolled system. Any vendor or provider promising guaranteed placement in AI answers should be treated skeptically; this claim is explicitly and repeatedly contradicted across the research reviewed here.
- Citation logic differs meaningfully by platform and changes frequently as underlying models and retrieval systems are updated, so tactics validated in early 2026 may lose effectiveness without changing anything about the content itself — this is a moving target, not a stable ranking algorithm.
- Because different engines draw on different source pools, a strategy optimized for one platform's citation logic (for example, Copilot's reliance on LinkedIn for B2B queries) may transfer poorly to another (for example, Claude's reported preference for long-form depth).
- Measurement infrastructure is immature. Attribution of a conversion to an AI-assistant referral is harder to verify than a standard UTM-tagged organic click, and industry-reported conversion-rate figures for AI-referred traffic vary widely by source and methodology.

9. Conclusion

Answer Engine Optimization in 2026 is best understood not as a replacement for SEO but as an additional, answer-specific layer built on top of it: a page still needs to be crawlable, indexed, and topically authoritative to enter an engine's candidate pool, and it additionally needs to be structured, evidenced, and kept current enough to survive the synthesis and citation stage. The strongest empirical evidence — from the founding Princeton/Georgia Tech study through its 2025–2026 replications — supports a fairly narrow, well-defined set of tactics: substantiate claims with statistics and citations, structure content for passage-level extraction, and keep high-intent content fresh. Broader claims common in marketing material, including guarantees of AI placement or one-size-fits-all tactics that work identically across every assistant, are not supported by the underlying research and should be treated

with the same skepticism a careful reader would apply to any early-stage, fast-moving optimization discipline.

References

Aggarwal, P., Murahari, V., Rajpurohit, T., Kalyan, A., Narasimhan, K., & Deshpande, A. (2024). Generative Engine Optimization (GEO): Improving Visibility in AI-Generated Search. Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD 2024). arxiv.org/abs/2311.09735

Adobe Analytics. (2026). AI-Referred Traffic to U.S. Retail Sites: Q1 2026 Analysis. Adobe Digital Insights, April 2026. business.adobe.com/blog/ai-traffic-surge-retail-sites-not-machine-readable

Ahrefs (Law, R., & Guan, X.). (2026). Update: AI Overviews Reduce Clicks by 58%. ahrefs.com/blog/ai-overviews-reduce-clicks-update

AirOps. (2026). State of AI Search Report 2026: Citation Patterns Across ChatGPT, Perplexity, Google AI Overview, and Gemini. airops.com/blog/aeo-answer-engine-optimization

Amsive. (2025). AI Overviews and Organic CTR: An Analysis of 700,000 Keywords Across 10 Websites and 5 Industries. amsive.com/insights/seo/google-ai-overviews-new-research-reveals-how-to-navigate-click-drop-off

Astiva AI. (2026). How to Optimize Content for AI Citations: 2026 LLM Guide. astiva.ai/blog/optimize-content-ai-citations-llm

AI Thinker Lab. (2026). Generative Engine Optimization (GEO) 2026: Princeton-Backed Playbook for AI Search. aithinkerlab.com/generative-engine-optimization-2026

Blck Alpaca. (2026). The Princeton GEO Study: Methodology, Results and Critique. blckalpaca.at/en/knowledge-base/seo-geo/geo-generative-engine-optimization

Conductor. (2026). AI Referral Traffic Benchmarks 2026. conductor.com/academy/aeo-geo-benchmarks-report

CXL. (2026). Answer Engine Optimization (AEO): The Comprehensive Guide for 2026. cxl.com/blog/answer-engine-optimization-aeo-the-comprehensive-guide

5W Public Relations. (2026). The State of AI Citations 2026: A Synthesis of 680 Million Tracked AI Citations Across ChatGPT, Claude, Perplexity, Gemini, and Google AI Overviews. Press release, June 2, 2026. prnewswire.com/news-releases/chatgpts-new-gatekeepers-wikipedia-reddit-and-the-sites-now-shaping-what-america-buys-302778843.html

DerivateX. (2026). The Princeton GEO Paper in Plain English: 5 Tactics That Boost AI Citation by 40%. derivatex.agency/blog/princeton-geo-paper-plain-english

Evergreen Media. (2026). Answer Engine Optimization (AEO): AI Visibility in 2026. evergreen.media/en/guide/answer-engine-optimization

Frase.io. (2026). Answer Engine Optimization: Complete AEO Guide. frase.io/blog/what-is-answer-engine-optimization-the-complete-guide-to-getting-cited-by-ai

Gartner. (2024). Gartner Predicts Search Engine Volume Will Drop 25% by 2026, Due to AI Chatbots and Other Virtual Agents. Press release, February 19, 2024. gartner.com/en/newsroom/press-releases/2024-02-19-gartner-predicts-search-engine-volume-will-drop-25-percent-by-2026-due-to-ai-chatbots-and-other-virtual-agents

HubSpot. (2026). Answer Engine Optimization Trends in 2026: How AEO Is Transforming the Landscape. blog.hubspot.com/marketing/answer-engine-optimization-trends

Mekaa. (2026). GEO: What Is Generative Engine Optimization? 2026 Guide. mekaa.co/en/blog/geo-what-is-generative-engine-optimization-2026-complete-guide

Percepture. (2026). Answer Engine Optimization AEO Strategies 2026. percepture.com/geo-insights/answer-engine-optimization-aeo-strategies-2026

Pew Research Center. (2025). Google Users Are Less Likely to Click on Links When an AI Summary Appears in the Results. July 22, 2025. pewresearch.org/short-reads/2025/07/22/google-users-are-less-likely-to-click-on-links-when-an-ai-summary-appears-in-the-results

Profound. (2025). AI Platform Citation Patterns: How ChatGPT, Google AI Overviews, and Perplexity Source Information. tryprofound.com/blog/ai-platform-citation-patterns

Position Digital. (2026). Answer Engine Optimization (AEO): 6 Best Practices for 2026. position.digital/blog/answer-engine-optimization-best-practices

Puerto, et al. (2025). C-SEO Bench: A Benchmark for Conversational SEO Tactics. arxiv.org/abs/2506.11097

rygr. (2026). Why AEO (Answer Engine Optimization) Is Critical to 2026 Marketing Planning. rygr.us/2026/02/04

Seer Interactive. (2025). AIO Impact on Google CTR: September 2025 Update. seerinteractive.com/insights/aio-impact-on-google-ctr-september-2025-update

Semrush. (2025). The Most-Cited Domains in AI: A 3-Month Study. semrush.com/blog/most-cited-domains-ai

Similarweb. (2025). Zero-Click Search Trends, May 2024–May 2025. As reported by Search Engine Journal: searchenginejournal.com/impact-of-ai-overviews-how-publishers-need-to-adapt

Surmado. (2026). Answer Engine Optimization: The Complete AEO and GEO Guide for 2026. surmado.com/blog/answer-engine-optimization-aeo-geo-guide

The Digital Bloom. (2026). 2025 AI Visibility Report: How LLMs Choose What Sources to Mention. thedigitalbloom.com/learn/2025-ai-citation-llm-visibility-report

Wellows. (2025). LLM Citation Trends That Matter in AI Search. wellows.com/blog/llm-citation-trends-for-ai-search

ALM Corp. (2026). Answer Engine Optimization (AEO) in 2026: Practical Playbook. almcorp.com/blog/answer-engine-optimization-2026

Various authors. (2026). Generative Engine Optimization at Scale: Measuring Brand Visibility Across AI Search Engines. arxiv.org/pdf/2606.20065

marvn.ai. (2026). marvn: The AI Answer Engine for iGaming. marvn.ai